

# BRIDGING THE ENSEMBLE KALMAN AND PARTICLE FILTER

MARCO FREI AND HANS R. KÜNSCH<sup>1</sup>

**ABSTRACT.** In many applications of Monte Carlo nonlinear filtering, the propagation step is computationally expensive, and hence, the sample size is limited. With small sample sizes, the update step becomes crucial. Particle filtering suffers from the well-known problem of sample degeneracy. Ensemble Kalman filtering avoids this, at the expense of treating non-Gaussian features of the forecast distribution incorrectly. Here we introduce a procedure which makes a continuous transition indexed by  $\gamma \in [0, 1]$  between the ensemble and the particle filter update. We propose automatic choices of the parameter  $\gamma$  such that the update stays as close as possible to the particle filter update subject to avoiding degeneracy. In various examples, we show that this procedure leads to updates which are able to handle non-Gaussian features of the prediction sample even in high-dimensional situations.

## 1. INTRODUCTION

State space models consist of a (discrete or continuous time) Markov process which is partially observed at discrete time points and subject to independent random errors. Estimation of the state at time  $t$  given observations up to the same time is called filtering or data assimilation. Since exact computations are possible essentially only in linear Gaussian situations, mostly Monte Carlo methods are used for filtering. In many environmental applications, in particular in atmospheric physics, oceanography and reservoir modelling, the dimension of the state is, however, very large and the computational costs to propagate the state forward in time are huge, which severely limits the potential sample size for Monte Carlo filtering methods. Particle filters (Gordon et al., 1993; Pitt and Shephard, 1999; Doucet et al., 2000) suffer from the well-known problem of sample degeneracy (Snyder et al., 2008). In contrast, the ensemble Kalman filter (Evensen, 1994; Burgers et al., 1998; Houtekamer and Mitchell, 1998) can handle some problems where the dimensions of states and observations are large, and the number of replicates is small, but at the expense of incorrectly treating non-Gaussian features of the forecast distribution that arise in nonlinear systems.

To relax the Gaussian assumption, two paradigms are predominant: mixture filters that approximate the forecast distribution as a mixture of Gaussians (Bengtsson et al., 2003; Sun et al., 2009; Døvera and Della Rossa, 2011; Frei and Künsch, 2011; Hoteit et al., 2012; Rezaie and Eidsvik, 2012), and sequential importance samplers that use the ensemble Kalman filter as a proposal distribution (Mandel and Beezley, 2009; Papadakis et al., 2010). In this article, we introduce an update scheme that blends these two flavours: A Gaussian mixture proposal obtained from an ensemble Kalman filter update based on a tempered likelihood is corrected by a particle filter update. In this

---

<sup>1</sup>SEMINAR FOR STATISTICS, ETH ZÜRICH, CH-8092 ZÜRICH, SWITZERLAND  
E-mail addresses: frei@stat.math.ethz.ch, kuensch@stat.math.ethz.ch.  
Date: July 31, 2012.

way we do not have to fit a Gaussian mixture to the forecast sample nor do we have to approximate the ratio of the predictive density to the proposal density. A further advantage of our procedure is that we can implement these two steps in such a way that the particle weights do not depend on artificial observation noise variables and the resampling avoids ties.

Our procedure depends on a single tuning parameter  $\gamma \in [0, 1]$ , which allows continuous interpolation between the ensemble Kalman filter ( $\gamma = 1$ ) and the particle filter ( $\gamma = 0$ ). Hence, the parameter  $\gamma$  controls the bias-variance trade-off between a correct update and maintaining the diversity of the sample. It can be chosen without prior knowledge based on a suitable measure of diversity like effective sample size ESS (Liu, 1996), or the expected number of Gaussian components which are represented in the resample.

The rest of the article is organized as follows. In Section 2, we detail the problem setting, introduce some notation, and provide background material. In Section 3, we present our new method and discuss implementational aspects. In Section 4, we discuss the choice of the tuning parameter  $\gamma$ . In Section 5, we consider numerical examples that involve single updates only; based on different prior specifications, we examine the differences of our method in comparison to the ensemble Kalman filter and the particle filter. In Section 6, we consider examples that involve many update cycles in two common test beds. Section 7 contains an outlook to possible generalizations.

## 2. PROBLEM SETTING, NOTATION AND BACKGROUND MATERIAL

We consider a dynamical system with state variable  $(x_t \in \mathbb{R}^q : t = 0, 1, \dots)$  and observations  $(y_t \in \mathbb{R}^r : t = 1, 2, \dots)$ . The state follows a deterministic or stochastic Markovian evolution, that is  $x_t = g(x_{t-1}, \xi_t)$  where the system noise  $\xi_t$  is independent of all past values  $x_s$  and all  $\xi_s$ ,  $s < t$ . There is no need to know the function  $g$  in explicit form, we only assume that for given  $x_{t-1}$  we are able to simulate from the distribution of  $g(x_{t-1}, \xi_t)$ . In particular, the evolution can be in continuous time, given by an ordinary or stochastic differential equation.

In all cases we assume linear observations with Gaussian noise:  $y_t = Hx_t + \epsilon_t$ ,  $\epsilon_t \sim \mathcal{N}(0, R)$ . This means that the likelihood for the state  $x_t$  given the observation  $y_t$  is  $\ell(x_t|y_t) = \varphi(y_t; Hx_t, R)$ . Here and in the following  $\varphi(x; \mu, \Sigma)$  denotes the (in general multivariate) normal density with mean  $\mu$  and covariance  $\Sigma$  at  $x$ . In the final section, we will discuss briefly how to adjust the method for non-Gaussian likelihoods.

We denote all observations up to time  $t$ ,  $(y_1, \dots, y_t)$  by  $y_{1:t}$ . The forecast distribution  $\pi_t^p$  at time  $t$  is the conditional distribution of  $x_t$  given  $y_{1:t-1}$ , and the filter distribution  $\pi_t^u$  at time  $t$  is the conditional distribution of  $x_t$  given  $y_{1:t}$ . In principle, these distributions can be computed recursively, alternating between propagation and update steps. The propagation step leads from  $\pi_{t-1}^u$  to  $\pi_t^p$ :  $\pi_t^p$  is the distribution of  $g(x_{t-1}, \xi_t)$  where  $x_{t-1} \sim \pi_{t-1}^u$  and  $\xi_t$  is independent of  $x_{t-1}$  and has the distribution given by the evolution of the system. The update step leads from  $\pi_t^p$  to  $\pi_t^u$  and is nothing else than Bayes formula:  $\pi_t^u(dx_t) \propto \ell(x_t|y_t)\pi_t^p(dx_t)$ . However, analytical computations are possible (essentially) only if the system evolution is also linear with additive Gaussian noise. Hence one resorts to Monte Carlo approximations, i.e., one represents  $\pi_t^p$  and  $\pi_t^u$  by ensembles (samples)  $(x_{t,j}^p)$  and  $(x_{t,j}^u)$  respectively. The members of these ensembles are called particles.

The propagation step just lets the particles evolve according to the dynamics of the state, that is we simulate according to the time evolution starting at  $x_{t-1,j}^u$  at time  $t-1$ ,  $x_{t,j}^p = g(x_{t-1,j}^u, \xi_{t,j})$ . However, the computational complexity of this step limits the number of particles, that is the size of the sample.

The (bootstrap) particle filter (Gordon et al., 1993) updates the forecast particles by weighting with weights proportional to the likelihood  $\ell(x_t|y_t)$  and converts this into an unweighted sample by resampling, i.e.,  $(x_{t,j}^u)$  is obtained by sampling from

$$(1) \quad \sum_{j=1}^N \omega_{t,j} \Delta_{x_{t,j}^p}, \quad \omega_{t,j} = \frac{\ell(x_{t,j}^p|y_t)}{\sum_{k=1}^N \ell(x_{t,k}^p|y_t)}.$$

Thus some of the forecast particles disappear and others are replicated. If the likelihood is quite peaked, which is the case in high dimensions with many independent observations, the weights will be heavily unbalanced, and the filter sample eventually degenerates since it concentrates on a single or a few particles; see Snyder et al. (2008). Auxiliary particle filters (Pitt and Shephard, 1999) can attenuate this behaviour to some extent, but they require good proposal distributions for the propagation, and an analytical expression for the transition densities.

The ensemble Kalman filter (Burgers et al., 1998; Houtekamer and Mitchell, 1998) makes an affine correction of the forecast particles based on the new observation  $y_t$  and artificial observation noise variables  $\epsilon_{t,j} \sim \mathcal{N}(0, R)$ :

$$x_{t,j}^u = x_{t,j}^p + K(\hat{P}_t^p)(y_t - Hx_{t,j}^p + \epsilon_{t,j})$$

where  $\hat{P}_t^p$  is an estimate of the forecast covariance at time  $t$ , typically a regularized version of the sample covariance of  $(x_{t,j}^p)$ , and  $K(P)$  is the Kalman gain  $K(P) = PH'(HPH' + R)^{-1}$ . This update formula is (asymptotically) correct under the assumption that the forecast distribution  $\pi_t^p$  is Gaussian (Le Gland et al., 2010). Although this is usually not valid, the update nevertheless has been found to work well in a variety of situations; see Evensen (2007) and references therein. For later use, we note that conditional on the forecast sample,

$$x_{t,j}^u \sim \mathcal{N}(x_{t,j}^p + K(\hat{P}_t^p)(y_t - Hx_{t,j}^p), K(\hat{P}_t^p)RK(\hat{P}_t^p)').$$

Therefore the filter sample can be considered as a balanced sample from the (conditional) Gaussian mixture

$$(2) \quad \frac{1}{N} \sum_{j=1}^N \mathcal{N}(x_{t,j}^p + K(\hat{P}_t^p)(y_t - Hx_{t,j}^p), K(\hat{P}_t^p)RK(\hat{P}_t^p)').$$

Here, “balanced sample” simply means that we draw exactly one realization from each of the  $N$  equally weighted Gaussian components.

### 3. A BRIDGE BETWEEN ENSEMBLE AND PARTICLE UPDATES: THE ENSEMBLE KALMAN PARTICLE FILTER

**3.1. The new method.** We consider here the update at a single fixed time  $t$  and thus suppress  $t$  in the notation. We follow the “progressive correction” idea (Musso et al., 2001) and write

$$\pi^u(dx) \propto \pi^{u,\gamma}(dx)\ell(x|y)^{1-\gamma}, \quad \pi^{u,\gamma}(dx) \propto \pi^p(dx)\ell(x|y)^\gamma$$

where  $0 \leq \gamma \leq 1$  is arbitrary. Our approach is to use an ensemble Kalman filter update to go from  $\pi^p$  to  $\pi^{u,\gamma}$ , and a particle filter update to go from  $\pi^{u,\gamma}$  to  $\pi^u$ . The rationale behind this two-stage procedure is to achieve a compromise between sample diversity and systematic error due to non-Gaussian features of  $\pi^p$ . The former is large if  $\gamma$  is close to one because the ensemble Kalman filter update draws the particles closer to the observation  $y$ , and the exponent  $1 - \gamma$  dampens the ratio of any two resampling probabilities. The latter is small if  $\gamma$  is small.

Since  $\ell(x|y)^\gamma \propto \varphi(y; Hx, R/\gamma)$ , and

$$(3) \quad PH'(HPH' + R/\gamma)^{-1} = \gamma PH'(\gamma HPH' + R)^{-1} = K(\gamma P)$$

the ensemble filter update is straightforward: We only need to compute the gain with the reduced covariance  $\gamma \hat{P}^p$ . The particle update then resamples with weights proportional to  $\varphi(y; Hx, \gamma R)$ . However, there are two immediate drawbacks to such an algorithm: the particle weights depend on the artificial observation noise variables which are needed for the ensemble Kalman filter update, and the resampling introduces tied values. We show next how to address both points. By (2), we can write

$$(4) \quad \pi^{u,\gamma} \approx \pi_{\text{EnKF}}^{u,\gamma} = \frac{1}{N} \sum_{j=1}^N \mathcal{N}(\nu_j^{u,\gamma}, Q(\gamma, \hat{P}^p))$$

where

$$(5) \quad \nu_j^{u,\gamma} = x_j^p + K(\gamma \hat{P}^p)(y - Hx_j^p),$$

$$(6) \quad Q(\gamma, \hat{P}^p) = \frac{1}{\gamma} K(\gamma \hat{P}^p) R K(\gamma \hat{P}^p)'.$$

Instead of sampling from (4) and applying a particle correction, we delay the ensemble Kalman filter sampling step, and update (4) analytically. This is easy because the update of a Gaussian mixture by a Gaussian likelihood is again a Gaussian mixture whose parameters can be computed easily (Alspach and Sorenson, 1972). We obtain

$$(7) \quad \pi^u \approx \pi_{\text{EnKPF}}^u = \sum_{j=1}^N \alpha_j^{u,\gamma} \mathcal{N}(\mu_j^{u,\gamma}, P^{u,\gamma})$$

where EnKPF stands for ensemble Kalman particle filter and

$$(8) \quad \alpha_j^{u,\gamma} \propto \varphi(y; H\nu_j^{u,\gamma}, HQ(\gamma, \hat{P}^p)H' + \frac{1}{1-\gamma}R),$$

$$(9) \quad \mu_j^{u,\gamma} = \nu_j^{u,\gamma} + K((1-\gamma)Q(\gamma, \hat{P}^p))(y - H\nu_j^{u,\gamma}),$$

$$(10) \quad P^{u,\gamma} = (I - K((1-\gamma)Q(\gamma, \hat{P}^p))H)Q(\gamma, \hat{P}^p).$$

The update consists now in sampling from (7). The mixture proportions  $\alpha_j^{u,\gamma}$  do not depend on the artificial observation noise variables, and even if one  $\alpha_j^{u,\gamma}$  dominates, there is still some diversity in the filter sample because the covariance  $P^{u,\gamma}$  is not zero if  $\gamma > 0$ .

Sampling from the  $j$ -th component of (7) can be done as follows: Let  $\epsilon_1$  and  $\epsilon_2$  be two independent  $\mathcal{N}(0, R)$  random variables. Then

$$(11) \quad x^{u,\gamma} = x_j^p + K(\gamma \hat{P}^p)(y + \frac{\epsilon_1}{\sqrt{\gamma}} - Hx_j^p) = \nu_j^{u,\gamma} + K(\gamma \hat{P}^p) \frac{\epsilon_1}{\sqrt{\gamma}}$$

clearly has distribution  $\mathcal{N}(\nu_j^{u,\gamma}, Q(\gamma, \hat{P}^p))$ , and thus by standard arguments

$$x^u = x^{u,\gamma} + K((1-\gamma)Q(\gamma, \hat{P}^p)) \left( y + \frac{\epsilon_2}{\sqrt{1-\gamma}} - Hx^{u,\gamma} \right)$$

is a sample from  $\mathcal{N}(\mu_j^{u,\gamma}, P^{u,\gamma})$ . Hence there is no need to compute a square root of  $P^{u,\gamma}$ .

To summarize, given a forecast ensemble  $(x_j^p)$  and an observation  $y$ , the ensemble Kalman particle filter consists of the following steps:

---

**Algorithm 1** Ensemble Kalman particle filter

---

1. Compute the estimated forecast covariance  $\hat{P}^p$ .
  2. Choose  $\gamma$  and compute  $K(\gamma \hat{P}^p)$  according to (3) and  $\nu_j^{u,\gamma}$  according to (5).
  3. Compute  $Q(\gamma, \hat{P}^p)$  according to (6) and the weights  $\alpha_j^{u,\gamma}$  according to (8).
  4. Choose indices  $I(j)$  by sampling from the weights  $\alpha_j^{u,\gamma}$  with some balanced sampling scheme; e.g., equation (12) in Künsch (2005).
  5. Generate  $\epsilon_{1,j} \sim \mathcal{N}(0, R)$  and set  $x_j^{u,\gamma} = \nu_{I(j)}^{u,\gamma} + K(\gamma \hat{P}^p) \frac{\epsilon_{1,j}}{\sqrt{\gamma}}$ .
  6. Compute  $K((1-\gamma)Q(\gamma, \hat{P}^p))$ , generate  $\epsilon_{2,j} \sim \mathcal{N}(0, R)$  and set  $x_j^u = x_j^{u,\gamma} + K((1-\gamma)Q(\gamma, \hat{P}^p)) \left( y + \frac{\epsilon_{2,j}}{\sqrt{1-\gamma}} - Hx_j^{u,\gamma} \right)$ .
- 

Because matrix inversion is continuous, it is easy to check that as  $\gamma \rightarrow 0$ ,  $\nu_j^{u,\gamma} \rightarrow x_j^p$ ,  $Q(\gamma, \hat{P}^p) \rightarrow 0$ ,  $\alpha_j^{u,\gamma} \rightarrow \varphi(y; Hx_j^p, R)$ ,  $\mu_j^{u,\gamma} \rightarrow x_j^p$  and  $P^{u,\gamma} \rightarrow 0$ . Hence in the limit  $\gamma \rightarrow 0$ , we obtain the particle filter update. Similarly, in the limit  $\gamma \rightarrow 1$  we obtain the ensemble Kalman filter update because for  $\gamma \rightarrow 1$   $(HQ(\gamma, \hat{P}^p)H' + R/(1-\gamma))^{-1}$  converges to zero and thus  $\alpha_j^{u,\gamma} \rightarrow 1/N$ . The ensemble Kalman particle filter therefore provides a continuous interpolation between the particle and the ensemble Kalman filter.

**3.2. Modifications in high dimensions.** If the dimension  $q$  of the state space is larger than the number  $N$  of particles, then the variability of the usual sample covariance is huge, and one should use a regularized version as estimate  $\hat{P}^p$ . In the context of the ensemble Kalman filter, the standard regularization technique is the use of a tapered estimate, that is we multiply the sample covariance matrix by a correlation matrix which is zero as soon as the distance between the two components of the state is larger than some threshold, see e.g., Houtekamer and Mitchell (2001); Furrer and Bengtsson (2007). If also the error covariance matrix  $R$  and the observation matrix  $H$  are sparse, then this tapering has the additional benefit that the computation of  $\nu_j^{u,\gamma}$  in step 2 and of  $x_j^{u,\gamma}$  in step 5 is much faster because we do not need to compute  $K(\gamma \hat{P}^p)$  for this. It is sufficient to solve  $2N$  equations of the form  $(\gamma HPH' + R)x = b$  which is fast for sparse matrices. However, this advantage is lost because for  $Q(\gamma, \hat{P}^p)$  we need  $K(\gamma \hat{P}^p)$ , which is in general a full matrix. We could multiply the gain matrix by another taper in order to facilitate the computation of  $Q(\gamma, \hat{P}^p)$  and to turn  $Q(\gamma, \hat{P}^p)$  into a sparse matrix. This would then make steps 4 and 6 in the algorithm above faster

because again all we need to do is to solve  $2N$  equations of the form  $((1-\gamma)HQ(\gamma, \hat{P}^p)H' + R)x = b$ . Alternatively, one could also replace the gain by the optimal matrix with a given sparsity pattern, i.e., using a localized update in grid space, see Sakov and Bertino (2010).

Because  $K(\gamma \hat{P}^p)$  is only used to compute  $Q(\gamma, \hat{P}^p)$ , a simpler alternative which avoids computing gain matrices is to generate the values  $K(\gamma \hat{P}^p)\epsilon_{1,j}/\sqrt{\gamma}$  needed in step 5 before step 3 and 4, and then to replace  $Q(\gamma, \hat{P}^p)$  by a sparse regularized version of the sample covariance of these values. If this approach is taken, it is usually feasible to generate more than  $N$  such values in order to reduce the Monte Carlo error in the regularized sample covariance matrix.

**3.3. Consistency of the ensemble Kalman particle filter in the Gaussian case.** We establish consistency of the ensemble Kalman particle filter for any  $\gamma$  as the ensemble size  $N$  tends to infinity, provided that the forecast sample is iid normal. We assume that all random quantities are defined on some given probability space  $(\Omega, \mathcal{F}, P)$  and “almost surely” is short for “P-almost surely”. The observation  $y$  is considered to be fixed (nonrandom). The superscript  $\cdot^N$  is added to any random quantity that depends on the ensemble size  $N$ . We use the following notion of convergence: A sequence  $(\pi^N)_{N \in \mathbb{N}}$  of random probability measures converges almost surely weakly to the probability measure  $\pi$  if  $\int h(x)\pi^N(dx)$  converges almost surely to  $\int h(x)\pi(dx)$  for any continuous and bounded function  $h$ .

**Theorem 1.** *Suppose that  $(x_j^p)_{j \in \mathbb{N}}$  is an iid sample from  $\pi^p = \mathcal{N}(\mu^p, P^p)$ . Then, for any  $\gamma \in [0, 1]$ , the sequence  $\pi_{\text{EnKPF}}^{u,N}$  as defined in (7) converges almost surely weakly to the true posterior  $\pi^u(dx) \propto \varphi(y; Hx, R)\pi^p(dx)$ . Additionally, if  $(x_j^{u,N})_{j=1,\dots,N}$  is a conditionally iid sample from  $\pi_{\text{EnKPF}}^{u,N}$ , then also  $\frac{1}{N} \sum_{j=1}^N \Delta_{x_j^{u,N}}$  converges almost surely weakly to  $\pi^u$ .*

A proof is given in the appendix. Notice that if balanced sampling is used to sample from  $\pi_{\text{EnKPF}}^u$ , the particles are no longer conditionally iid, and the arguments become more complicated; see Künsch (2005) for a discussion in the context of auxiliary particle filters. Inspection of the proof of Theorem 1 shows that if  $\pi^p$  is non-Gaussian with finite second moments, then  $\pi_{\text{EnKPF}}^u$  still converges almost surely weakly to a nonrandom limit distribution  $\pi_{\text{EnKPF}}^{u,\infty}$ . The limit distribution depends on  $\gamma$  and generally differs from the correct posterior  $\pi^u$  for  $\gamma > 0$ . The limit cannot be easily identified, and in particular it is difficult to quantify the systematic error as a function of  $\gamma$ . Using similar arguments as in Randles (1982), it is also possible to show that

$$N^{1/2} \left( \int h(x)\pi_{\text{EnKPF}}^{u,N}(dx) - \int h(x)\pi_{\text{EnKPF}}^{u,\infty}(dx) \right) \rightarrow \mathcal{N}(0, V)$$

weakly, where the asymptotic covariance  $V$  depends on  $h$ ,  $\pi^p$  and  $\gamma$ . In general,  $V$  is analytically intractable, and we cannot verify if  $V$  decreases as a function of  $\gamma$ , as we expect.

#### 4. CHOICE OF $\gamma$

**4.1. Asymptotics of weights.** Recall that for  $\gamma = 0$  the method is exactly the particle filter, and for  $\gamma = 1$  it is exactly the ensemble Kalman filter. Hence it is clear that there is a range of values  $\gamma$  where we obtain an interesting compromise between the two methods in the sense that the weights  $(\alpha_j^{u,\gamma})$  are neither uniform nor degenerate. We try to provide some theoretical insight where this range of values  $\gamma$  is, and later we develop a criterion which chooses a good value  $\gamma$  automatically.

We want to see how the weights  $\alpha_j^{u,\gamma}$  in (7) behave as a function of  $\gamma$  when the dimension of the observations is large. By definition

$$\begin{aligned}\alpha_j^{u,\gamma} &\propto \exp\left(-\frac{1}{2}(y - H\nu_j^{u,\gamma})'(HQ(\gamma, \hat{P}^p)H' + \frac{1}{1-\gamma}R)^{-1}(y - H\nu_j^{u,\gamma})\right) \\ &\propto \exp\left(-\frac{1}{2}(x_j^p - \mu^p)'\hat{C}_\gamma(x_j^p - \mu^p) + \hat{d}'_\gamma(x_j^p - \mu^p)\right)\end{aligned}$$

where  $\mu^p$  is the prediction mean,

$$\begin{aligned}\hat{C}_\gamma &= (1-\gamma)H'(I - \hat{K}_\gamma' H')((1-\gamma)H\hat{Q}_\gamma H' + R)^{-1}(I - H\hat{K}_\gamma)H, \\ \hat{d}_\gamma &= (1-\gamma)H'(I - \hat{K}_\gamma' H')((1-\gamma)H\hat{Q}_\gamma H' + R)^{-1}(I - H\hat{K}_\gamma)(y - H\mu^p)\end{aligned}$$

and  $\hat{K}_\gamma$  and  $\hat{Q}_\gamma$  stand for  $K(\gamma\hat{P}^p)$  and  $Q(\gamma, \hat{P}^p)$ .

The following lemma gives an approximate formula for the variance of the  $\alpha_j^{u,\gamma}$ .

**Lemma 1.** *Define approximate weights by*

$$\tilde{\alpha}_j^{u,\gamma} = \frac{1}{N} \frac{\exp(-\frac{1}{2}(x_j^p - \mu^p)'C_\gamma(x_j^p - \mu^p) + d'_\gamma(x_j^p - \mu^p))}{\mathbf{E}[\exp(-\frac{1}{2}(x_j^p - \mu^p)'C_\gamma(x_j^p - \mu^p) + d'_\gamma(x_j^p - \mu^p))]}.$$

where  $C_\gamma$  and  $d_\gamma$  are as defined above, but with the true forecast covariance  $P^p$  instead of  $\hat{P}^p$ . If the forecast sample is iid  $\mathcal{N}(\mu^p, P^p)$ , then

$$N^2 \text{Var} [\tilde{\alpha}_j^{u,\gamma}] = \frac{\det(P^p C_\gamma + I)}{\det(2P^p C_\gamma + I)^{1/2}} \exp(d'_\gamma((C_\gamma + (P^p)^{-1}/2)^{-1} - (C_\gamma + (P^p)^{-1})^{-1})d_\gamma) - 1.$$

As  $\gamma \uparrow 1$ , we have

$$(12) \quad N^2 \text{Var} [\tilde{\alpha}_j^{u,\gamma}] \sim (1-\gamma)^2 \left( \frac{1}{2} \text{tr}(HP^p H' M) + (y - H\mu^p)' M (y - H\mu^p) \right),$$

where  $M = (I - K_1' H')R^{-1}(I - HK_1)HP^p H'(I - K_1' H')R^{-1}(I - HK_1)$ .

A proof is given in the appendix. The matrix  $M$  is positive definite and also the trace in formula (12) is positive. Therefore we expect that the variance of the weights is of the order  $\mathcal{O}(N^{-2}(1-\gamma)^2 q)$ : This is true if  $P^p$ ,  $R$  and  $H$  are all multiples of the identity, and there is no reason for a different behavior in other cases. This suggests that for high-dimensional observations, we need to choose  $\gamma$  close to one in order to avoid degeneracy. Note however that the final update can still differ from the ensemble Kalman filter update, even if the largest part of the update occurs with the ensemble Kalman filter.

**4.2. Criteria for the selection of  $\gamma$ .** Because the Kalman filter update uses only the first two moments of the forecast distribution, it seems plausible that for non-Gaussian forecast distributions, the Kalman update will be less informative than the correct update. Hence, as long as the spread is a meaningful measure of uncertainty, we expect the Kalman update to have a larger spread than the correct update. This is not always true though: The variance of the correct posterior is only on average smaller than the variance  $(I - K(P^p)H)P^p$  of the ensemble Kalman filter update. In particular, this may fail to hold for some values of  $y$  if the prior is multimodal.

Still, this heuristic suggests that in some cases the spread of the update ensemble will increase monotonically in  $\gamma$  and we could choose  $\gamma$  such that the spread of the update is not smaller than a

factor  $\tau$  times the spread of an ensemble Kalman filter update (where  $\tau$  is maybe between 0.5 and 0.8). This means that we compare the variances of the Gaussian mixture  $\sum_{j=1}^n \frac{N_j}{N} \mathcal{N}(\mu_j^{u,\gamma}, P^{u,\gamma})$ , where  $N_j$  denotes the number of times component  $j$  in (7) has been selected, with the variances of the Kalman filter update, that is the diagonal elements of  $(I - K(\hat{P}^p)H)\hat{P}^p$ . However, this is computationally demanding, because we have to compute among other things  $P^{u,\gamma}$ ; compare the discussion in Section 3.2 above.

A simpler procedure is based on the standard deviations of the updated ensembles. Denote by  $\hat{\sigma}_i^u$  the standard deviation of the  $i$ -th component of the final update sample  $x_j^u$  (which depends on  $\gamma$ ) and by  $\hat{\sigma}_i^{u,En}$  the standard deviation of the update sample using the ensemble Kalman filter. Then we could choose the smallest  $\gamma$  such that  $\sum_i \hat{\sigma}_i^u \geq \tau \sum_i \hat{\sigma}_i^{u,En}$  or (in order to control not only the total spread but all marginal spreads) such that  $\sum_i \min(1, \hat{\sigma}_i^u (\hat{\sigma}_i^{u,En})^{-1}) \geq \tau N$ . This has the advantage that it is easy to compute also in high dimensions. The disadvantage is that it depends also on the generated noises. This can be reduced somehow by taking the same noises  $\epsilon_{1,j}$  and  $\epsilon_{2,j}$  for all values of  $\gamma$  under consideration.

A third possibility is to look only at the properties of the weights  $\alpha_j^{u,\gamma}$ . We can take as the measure of the sampling diversity the so-called effective sample size ESS (Liu, 1996), which is defined as

$$\text{ESS} = \frac{1}{\sum_j (\alpha_j^{u,\gamma})^2} = \frac{N}{1 + N \sum_j (\alpha_j^{u,\gamma} - 1/N)^2},$$

or the quantity

$$\text{DIV} = \sum_{j=1}^N \min(1, N \alpha_j^{u,\gamma}) = N(1 - \frac{1}{2} \sum_j |\alpha_j^{u,\gamma} - 1/N|),$$

which is the expected number of components that are chosen when generating  $(x_j^u)$  according to (7) with balanced sampling. Both measures are related to a distance between the  $\alpha_j^{u,\gamma}$  and uniform weights. Although there is no universal relation between the two criteria, in typical cases  $\text{ESS} < \text{DIV}$ , i.e., ESS is a more conservative measure of diversity. Both criteria do not take the spread in  $P^{u,\gamma}$  into account, which also increases if  $\gamma$  increases. Therefore, they give only a lower bound for the diversity, but they are easy to compute. We then choose  $\gamma$  as the smallest value for which  $\text{ESS} > \tau N$  or  $\text{DIV} > \tau N$ . In order to avoid excessive computations, we considered in the examples below only multiples of  $1/15$  as values for  $\gamma$  and used a binary search tree (with at most 4 search steps), assuming that the diversity is increasing in  $\gamma$ . We did not try to prove this assumption since the calculation is expected to be extremely tedious; the assumption is safe to make, though, since at the worst we end up with a too large  $\gamma$ . Alternatively one could use an approximation of ESS based on (12).

## 5. EXAMPLES OF SINGLE UPDATES

**5.1. Description of the setup.** We consider a single update for 4 situations. In all cases  $H = I$  and  $R = \sigma^2 I$ . There are 2 forecast samples  $(x_j^p)$  combined with 2 values  $y = y_1$  and  $y = y_2$  for the observations. The construction of the forecast sample starts with a sample  $z_j \sim \mathcal{N}_q(0, I)$ , which is then modified to introduce non-Gaussian features. More precisely, we consider the following situations:






FIGURE 1. Diversity  $N^{-1} \times \text{ESS}$  as a function of  $\gamma$ . Top row: Gaussian prior, bottom row: bimodal prior. Left column:  $y_1$ , right column:  $y_2$ . Dimension  $q$  of the state variable (from top line to bottom line):  $q = 10, 50, 250$ . The dotted lines show the approximate diversity computed from (12) with  $\mu^p$  and  $P^p$  estimated from the sample.

1. A Gaussian prior: We set  $\sigma = 0.5$ ,  $x_j^p = z_j$ ,  $y_1 = (0, \dots, 0)'$ , and  $y_2 = (1.5, 1.5, 0, \dots, 0)'$ . This means that the second observation contradicts the prior, although not excessively.
2. A bimodal prior: We set  $\sigma = 3$ ,  $x_j^p = z_j$  for  $j \leq N/2$  and  $x_j^p = z_j + (6, 0, \dots, 0)'$  for  $j > N/2$ ,  $y_1 = (-2, 0, \dots, 0)'$  and  $y_2 = (3, 0, \dots, 0)$ . In the former case, the true posterior is unimodal and in the latter case it is bimodal.

We take  $N = 50$  and  $q = 10, 50, 250$ . We generate one sample in dimension 250 and use the first  $q$  components. In all cases, we use a triangular taper with range 10, assuming that the states are values along a line (the optimal taper has range 0, since the true covariance  $P^p$  is diagonal). Without a taper, all procedures break down: They become overconfident as the dimension increases, and the estimated sampling diversity increases with  $q$ .

**5.2. Variation of  $\gamma$ .** We compute the update for  $\gamma \in \{0, 0.05, 0.10, \dots, 1\}$  and take as the measure of the sampling diversity the quantity  $N^{-1} \times \text{ESS}$  introduced above. Some results are shown in Figure 1. The diversity increases with  $\gamma$  and decreases with the dimension  $q$  as expected. In the bimodal case, even the particle filter does not degenerate, and typically small or moderate values of  $\gamma$  apparently give sufficient diversity.

**5.3. The updates of the first two coordinates.** We concentrate on the first two coordinates of  $x_j^{u,\gamma}$  which contain the non-Gaussian features (if present). We show the contours of the true update density (7) for the ensemble Kalman filter and for the filter with  $\gamma$  chosen such that the diversity  $\tau = N^{-1} \cdot \text{ESS}$  is approximately 40%. Figure 2 shows the results for bimodal prior with  $q = 250$ . In case of the Gaussian prior, the two plots (which are not shown here) are virtually identical. In the non-Gaussian situation, the combined filter is able to pick up some non-Gaussian features. In particular, the shape and not only the location depends on the observation, and the bimodality of the posterior is captured.

## 6. EXAMPLES OF FILTERING WITH MANY CYCLES

**6.1. The Lorenz 96 model.** The 40-variable configuration of the Lorenz 96 model (Lorenz and Emanuel, 1998) is governed by the ordinary differential equation

$$\frac{dX_t^k}{dt} = (X_t^{k+1} - X_t^{k-2})X_t^{k-1} - X_t^k + 8, \quad k = 1, \dots, 40.$$

where the boundary conditions are assumed to be cyclic, i.e.,  $X^k = X^{40+k}$ . The model is chaotic and mimics the time-evolution of a scalar meteorological quantity on a latitude circle. We adopt the same experimental setup as in Bengtsson et al. (2003) and Frei and Künsch (2011): Measurements of odd components  $X^{2k-1}$  with uncorrelated additive  $\mathcal{N}(0, 0.5)$  noise at observation times  $0.4 \times n$ ,  $n = 1, \dots, 2000$ , are taken. The large lead time produces a strongly nonlinear propagation step. The system is integrated using Euler's method with step size 0.001. Both the ensemble Kalman filter and ensemble Kalman particle filter are run with  $N = 400$  ensemble members. The true initial state and the initial ensemble members are randomly drawn from  $\mathcal{N}_{40}(0, I)$ . All sample covariance matrices are replaced by tapered estimates; for the sake of simplicity, we used the same taper matrix  $C$  throughout, namely the GC taper constructed from the correlation function given in (Gaspari and Cohn, 1999, equation (4.10)) with support half-length  $c = 10$ . For the ensemble Kalman particle filter, the parameter  $\gamma$  is chosen adaptively to ensure that the diversity  $\tau = N^{-1} \times \text{ESS}$  stays within a prespecified interval  $[\tau_0, \tau_1] \subset [0, 1]$  if possible. We also ran the filter proposed in Papadakis et al. (2010). For the given ensemble size, we were not able to obtain a non-divergent run; the filter collapsed after just a few cycles.

The filter performance is assessed via a scoring rule evaluated at observation times. Here, we use the root mean square error of the ensemble mean, and the continuous ranked probability score (Gneiting et al., 2007) for the first two state variables. More precisely, if  $X_t^k$  is the true solution at time  $t$ , and  $\hat{X}_t^k$  the mean of the updated ensemble, and  $\hat{F}_t^k(y)$  the marginal empirical cumulative distribution function of the updated ensemble, then the root mean square error of the ensemble mean at time  $t$  is

$$(13) \quad \text{RMSE}_t = \sqrt{\frac{1}{q} \sum_{k=1}^q (X_t^k - \hat{X}_t^k)^2}$$

and the continuous ranked probability score for the  $k$ th variable at time  $t$  is

$$(14) \quad \text{CRPS}_t^k = \int_{\mathbb{R}} \left( \hat{F}_t^k(y) - 1_{\{y \geq X_t^k\}} \right)^2 dy, \quad k = 1, 2$$

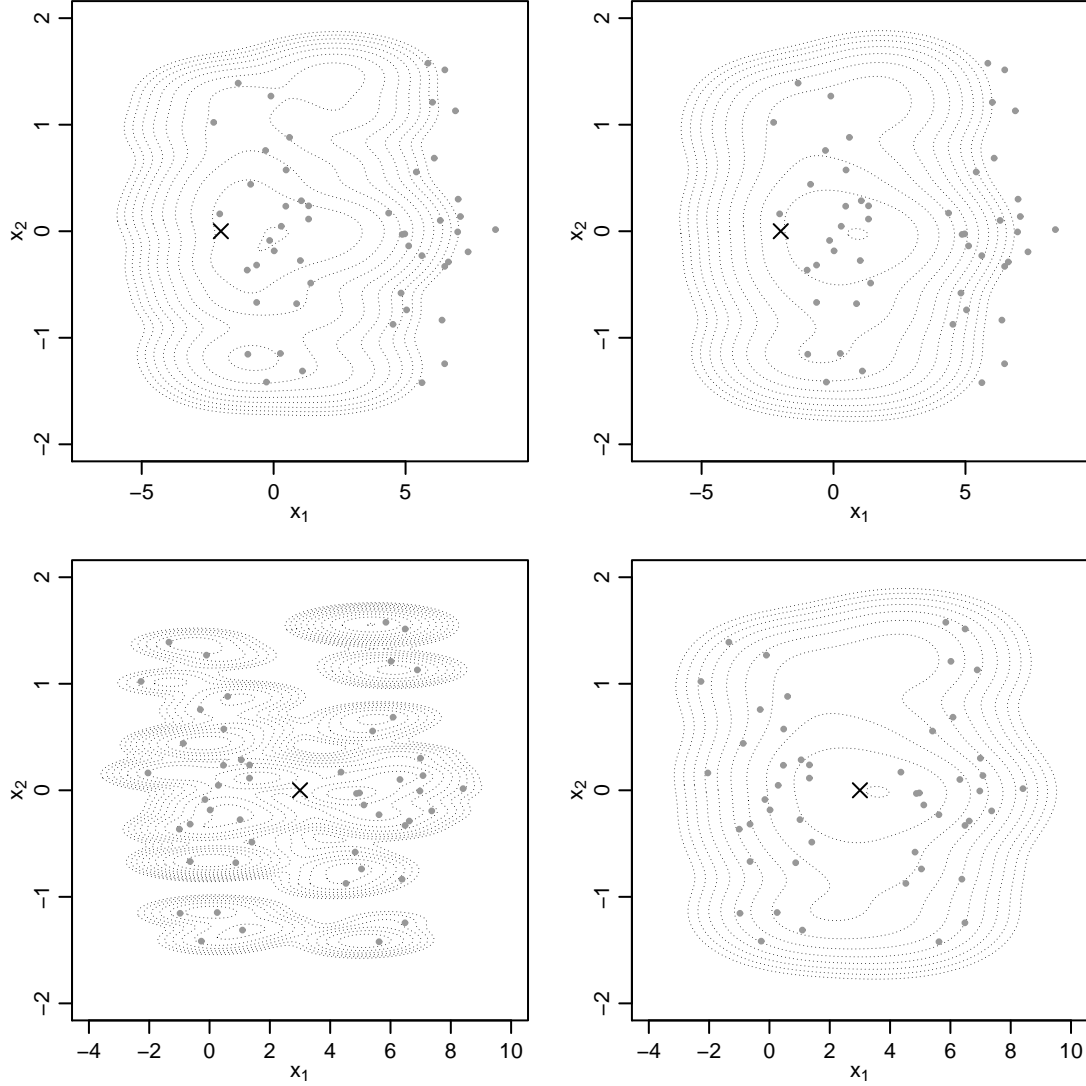


FIGURE 2. First two components of the update of the bimodal prior with  $q = 250$ . Upper row:  $y_1$ , lower row:  $y_2$ . Left column:  $\gamma$  chosen to achieve a diversity of about 40%. Right column: Ensemble Kalman filter. The prior sample is shown light grey, the observation is marked with a cross. The contours show the Gaussian mixture (7): Levels are equidistant on a log scale such that the lowest level corresponds to 1% of the maximum.

where  $t = 0.4 \times n, n = 1, \dots, 2000$ . Notice that for reasons of symmetry, we only consider the CRPS of the first two state variables (i.e., one observed and one unobserved variable).

Tables 1 and 2 compile summaries (first and ninth decile, mean and median) of the 2000 RMSE and CRPS values.

The gain over the ensemble Kalman filter achieved by the ensemble Kalman particle filter is substantial. In particular, as the CRPS values show, the ensemble Kalman particle filter is able to track the unobserved states much more accurately than the ensemble Kalman filter. Overall, the

	$[\tau_0, \tau_1]$	10%	50%	mean	90%
EnKF		0.56	0.81	0.87	1.25
EnKPF	[0.80, 0.90]	0.52	0.75	0.83	1.21
EnKPF	[0.50, 0.80]	0.51	0.73	0.80	1.18
EnKPF	[0.30, 0.60]	0.50	0.71	0.79	1.17
EnKPF	[0.25, 0.50]	0.49	0.70	0.78	1.16
EnKPF	[0.10, 0.30]	0.49	0.71	0.79	1.17

TABLE 1. Lorenz 96 system, experimental setup as given in Section 6.1: summary statistics of RMSE (13) over 2000 cycles using the ensemble Kalman filter (EnKF) and the ensemble Kalman particle filter (EnKPF) with constrained diversity  $\tau = N^{-1}\text{ESS} \in [\tau_0, \tau_1]$  for the weights.

	$[\tau_0, \tau_1]$	$X^1$				$X^2$			
		10%	50%	mean	90%	10%	50%	mean	90%
EnKF		0.12	0.22	0.32	0.65	0.14	0.38	0.57	1.18
EnKPF	[0.80, 0.90]	0.11	0.21	0.30	0.62	0.13	0.33	0.54	1.13
EnKPF	[0.50, 0.80]	0.11	0.21	0.29	0.61	0.12	0.32	0.51	1.10
EnKPF	[0.30, 0.60]	0.11	0.20	0.29	0.59	0.12	0.32	0.49	1.02
EnKPF	[0.25, 0.50]	0.10	0.20	0.28	0.58	0.11	0.31	0.48	1.00
EnKPF	[0.10, 0.30]	0.10	0.21	0.29	0.59	0.11	0.31	0.50	1.05

TABLE 2. Lorenz 96 system, experimental setup as given in Section 6.1: summary statistics of CRPS (14) over 2000 cycles for the state variables  $X^1$  (observed) and  $X^2$  (unobserved) using the ensemble Kalman filter (EnKF) and the ensemble Kalman particle filter (EnKPF) with constrained diversity  $\tau = N^{-1}\text{ESS} \in [\tau_0, \tau_1]$  for the weights.

results for the ensemble Kalman particle filter are comparable with those reported for the XEnKF in Frei and Künsch (2011). Arguably, the best performance of the ensemble Kalman particle filter is achieved with diversity constrained to  $[0.25, 0.50]$ , but the scores are surprisingly robust. Finally, we note that for smaller ensemble sizes, e.g.,  $N = 100$ , the ensemble Kalman particle filter still improves over the ensemble Kalman filter, but the results (which are not shown here) are less impressive. Nevertheless, one should keep in mind that with a particle filter, even with judicious tuning, much more particles are required to compete with the ensemble Kalman filter, as illustrated in Bocquet et al. (2010) (for a slightly different configuration of the Lorenz model).

**6.2. The Korteweg-de Vries equation.** We consider the Korteweg-de Vries equation on the circle (Drazin and Johnson, 1989):

$$\partial_t x + \partial_s^3 x + 3\partial_s x^2 = 0$$

with domain  $(s, t) \in [-1, 1) \times [0, \infty)$  and periodic boundary conditions,  $x(s = -1, t) = x(s = 1, t)$ . Versions of this equation have been used as test beds for data assimilation in, e.g., van Leeuwen (2003), Lawson and Hansen (2005), or Zupanski and Zupanski (2006). The spatial domain  $[-1, 1)$  is discretized using an equispaced grid with  $q = 128$  grid points. The spectral split step method is used to solve the equation numerically (with an explicit 4th order Runge-Kutta time step for the




FIGURE 3. Korteweg-de Vries equation. The left figure shows the initial 16-member ensemble (grey) at time  $t = 0$  together with the true solution (black). The right figure shows the predictive ensemble (grey) at the first observation time  $t = 0.01$  together with the observations (black bullets) and the true solution (black).

nonlinear part of the equation). As initial prior we take the random field

$$X(s, t = 0) = \exp\left(-\frac{s^2}{\eta^2}\right), \quad \log(\eta) \sim \mathcal{U}(\log(0.05), \log(0.3)).$$

For the truth, we use  $\eta = 0.2$ , and the initial ensemble is a (quasi-random) sample from  $X(s, t = 0)$ . The ensemble size is  $N = 16$ , and thus  $N \ll q$ . Six irregularly spaced observations with uncorrelated additive  $\mathcal{N}(0, 0.02)$  noise at observation times  $0.01 \times n$ ,  $n = 1, \dots, 10$ , are taken. For illustration, Figure 3 displays the initial 16-member ensemble and the predictive ensemble at the first observation time (together with the observations).

The particle filter, the ensemble Kalman filter and the ensemble Kalman particle filter are run (with no tapering applied). For the ensemble Kalman particle filter, we fix  $\gamma = 0.05$ , which ensures that  $\tau = N^{-1} \cdot \text{ESS}$  lies roughly in the interval  $[0.80, 0.90]$ . Since the particle filter degenerates very quickly for such a small ensemble, a benchmark run with  $N = 256$  particles is carried out. Figure 4 displays ensemble deviations from the true solution after 1 and 10 update cycles. Apparently, both the ensemble Kalman filter and ensemble Kalman particle filter track the true solution reasonably well, and the state uncertainty is well represented. In terms of error of the ensemble mean, there is not much difference between the ensemble Kalman filter and the ensemble Kalman particle filter. However, the ensemble Kalman particle filter produces particles that exhibit less dynamical inconsistencies. In Figure 5, for each filter, the particle with the most curvature after 10 update cycles is shown, where the curvature of a solution  $x(s, t)$  is defined by  $\int_{-1}^1 |\partial_s^2 x| (1 + (\partial_s x)^2)^{-3/2} ds$ , and a finite difference approximation is used for the discretized solutions. For reference, we note that the true solution (not shown in the plots) is virtually identical to the particle shown in the rightmost plot. The particle filter (which conserves dynamical constraints) yields very smooth particles, whereas the ensemble Kalman filter may produce wiggly, “unphysical” particles. The ensemble Kalman particle filter lies in-between: some particles still show slight dynamical imbalances, but these are much less pronounced than for the ensemble Kalman filter. Also, if a taper is applied to the forecast covariance

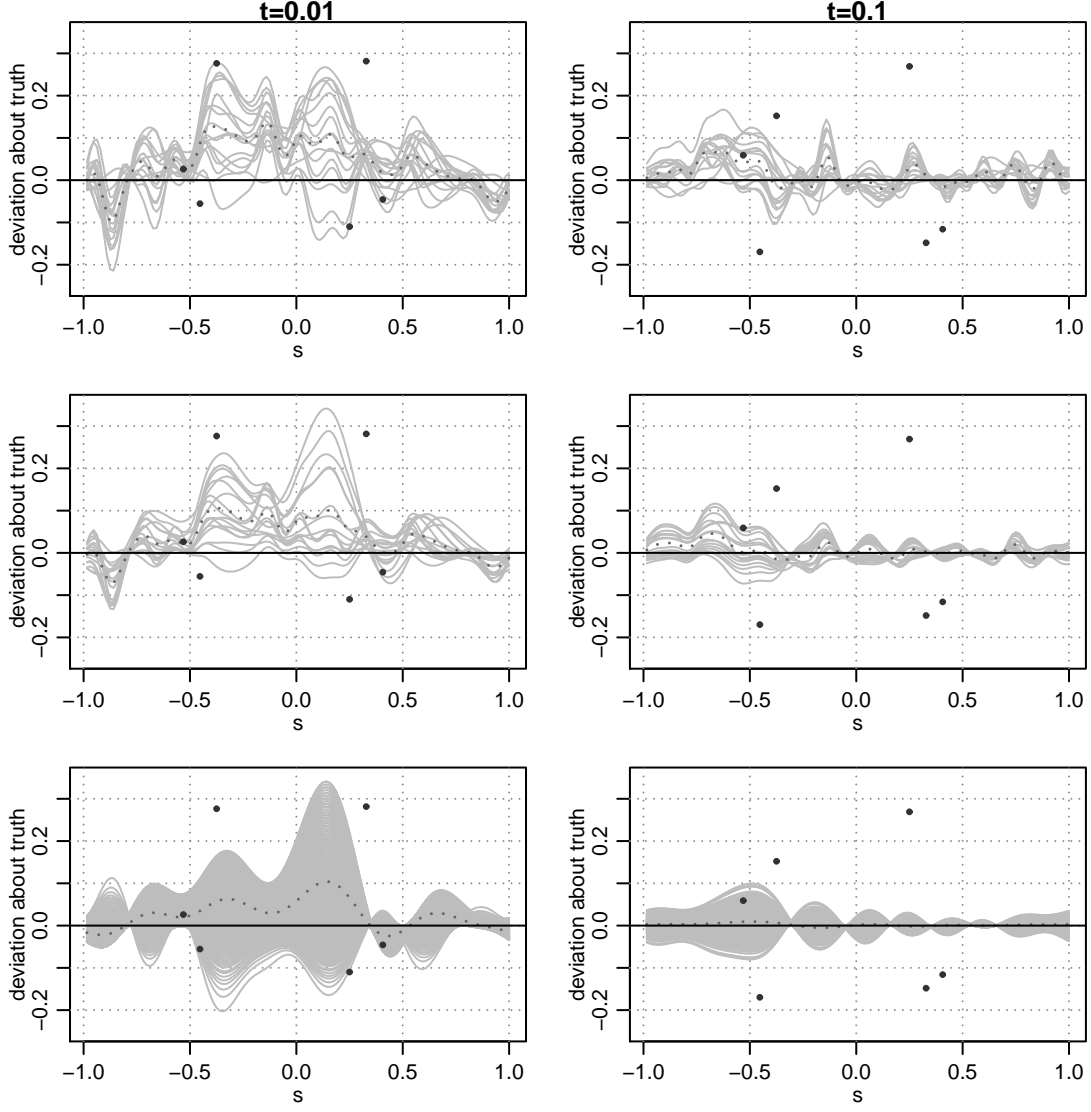


FIGURE 4. Korteweg-de Vries equation. Ensemble deviations about the truth, i.e., filtered ensemble minus true solution, after the 1st ( $t = 0.01$ , left panel) and 10th ( $t = 0.1$ , right panel) update cycle, for the ensemble Kalman filter and ensemble Kalman particle filter with  $N = 16$  particles (top two rows), and for the particle filter with  $N = 256$  particles (bottom row). The solid grey lines are the deviations, the dotted grey lines the average of the deviations, and the black bullets are the observations minus the truth.

matrix (which is not the case in the example here), the ensemble Kalman filter suffers even more from these imbalances.

## 7. OUTLOOK TO POSSIBLE GENERALIZATIONS

In the spirit of the “progressive correction” idea (Musso et al., 2001), the ensemble Kalman particle filter update could also be split up in several steps. We fix constants  $\gamma_i > 0$  and  $\delta_i > 0$  with  $\sum_{i=1}^L \gamma_i + \delta_i = 1$ . Then, for  $i = 1, \dots, L$ , we apply an ensemble Kalman filter update with likelihood




FIGURE 5. Korteweg-de Vries equation. Particle with most curvature after 10 update cycles ( $t = 0.1$ ), for the ensemble Kalman filter (left), ensemble Kalman particle filter (middle) and particle filter (right).

$\ell(x|y)^{\gamma_i}$ , followed by a particle filter update with likelihood  $\ell(x|y)^{\delta_i}$ , followed by the resampling step. It is only necessary to estimate the predictive covariance for the first step; for the subsequent steps,  $i = 2, \dots, L$ , we can compute the covariance analytically from the mixture representation (7) (for large  $q$ , this is numerically delicate, but the same remedies as discussed in Section 3.2 can be applied). We expect that the bias of such an iterative ensemble Kalman particle filter update is similar as for a single ensemble Kalman particle filter update with  $\gamma = \sum_{i=1}^L \gamma_i$ , but the variance will decrease with increasing  $L$  since the likelihoods become flatter. In the limiting case  $\sum_{i=1}^L \gamma_i \rightarrow 0$ , which corresponds to a full particle filter update, we conjecture that  $L = \mathcal{O}(q)$  is sufficient to retain the sampling diversity. This claim is supported by Beskos et al. (2012), who analyze the “tempering” idea in simpler but related situations.

A potential drawback of the ensemble Kalman particle filter (in comparison to non-Gaussian ensemble filters akin to Lei and Bickel (2012)) is its restriction to Gaussian linear observations. However, the idea of combining an ensemble Kalman filter and a particle filter update could also be used for arbitrary observation densities. Let  $H$  be a matrix that selects those components of the state variable that influence the observation, and assume that we have an approximation of the likelihood of the form  $\ell(Hx|y) \approx \varphi(g(y); Hx, R(y))$ . Then we can use this approximation for an ensemble Kalman filter update, and correct by a particle filter update with weights proportional to  $\ell(Hx_j^u|y)(\varphi(g(y); Hx_j^u, R(y)))^{-1}$ . In order to construct an approximation of the likelihood of the above form, we can use a Taylor approximation

$$\log \ell(Hx|y) \approx \log \ell(H\mu^p|y) + a(y)'H(x - \mu^p) + \frac{1}{2}(x - \mu^p)'H'b(y)H(x - \mu^p)$$

where  $a(y)$  and  $b(y)$  are the gradient and the Hessian, respectively, of the log likelihood. Then  $R(y) = -b(y)^{-1}$  and  $g(y) = R(y)a(y)$ . Alternatively, one could center the expansion at the mode of likelihood. Such an approximation is expected to work well in cases where the likelihood is log-concave, e.g., when  $y$  given  $x$  is Poisson with parameter  $\exp(x)$ .

#### ACKNOWLEDGMENTS

The authors thank Jo Eidsvik for fruitful discussions.

## APPENDIX

**Proof of Theorem 1.** Convergence of  $\pi_{\text{EnKPF}}^{u,N}$  to  $\pi^u$  implies convergence of  $\frac{1}{N} \sum_{j=1}^N \Delta_{x_j^{u,N}}$  to  $\pi^u$ , see Lemma 7 in Frei and Künsch (2011). It remains to establish the former convergence. To begin with, we introduce some notation used throughout the remainder. For ease of legibility, the dependence on  $N$  is dropped. An overbar  $\bar{\cdot}$  is added to any quantity to denote its population counterpart, in which  $\hat{P}^p$  has been replaced by  $P^p$ . For a measure  $\pi$  and a function  $h$ , we write  $\pi h = \int h(x) \pi(dx)$ . Expressions of the form  $A \rightarrow B$  are shorthand for  $A$  converges to  $B$  almost surely as  $N$  goes to  $\infty$ . Straightforward application of the strong law of large numbers shows that the population version  $\bar{\pi}_{\text{EnKPF}}^{u,N}$  converges to some nonrandom limit, and it is clear by construction that this limit equals  $\pi^u$ . Hence, it remains to prove that the population version of the ensemble Kalman particle filter has the same limit as the ensemble Kalman particle filter, i.e., we need to prove that

$$(15) \quad |\pi_{\text{EnKPF}}^u h - \bar{\pi}_{\text{EnKPF}}^u h| \rightarrow 0$$

for any continuous and bounded function  $h$ . In addition, we may assume that  $h$  is compactly supported on, say,  $S_h \subset \mathbb{R}^q$ , since these functions are still convergence determining for the weak topology. Write  $\|h\|_\infty = \max_x |h(x)|$ .

Observe that

$$(16) \quad \begin{aligned} |\pi_{\text{EnKPF}}^u h - \bar{\pi}_{\text{EnKPF}}^u h| &\leq \left| \sum_{j=1}^N \alpha_j^u \mathcal{N}(\mu_j^u, P^u) h - \sum_{j=1}^N \bar{\alpha}_j^u \mathcal{N}(\mu_j^u, P^u) h \right| \\ &\quad + \left| \sum_{j=1}^N \bar{\alpha}_j^u \mathcal{N}(\mu_j^u, P^u) h - \sum_{j=1}^N \bar{\alpha}_j^u \mathcal{N}(\bar{\mu}_j^u, \bar{P}^u) h \right|. \end{aligned}$$

In the following, we show that both terms on the right-hand side of (16) converge to 0, which proves (15). For later use, we note that  $\hat{P}^p \rightarrow P^p$ , and hence by continuity,

$$(17) \quad K(\gamma \hat{P}^p) \rightarrow K(\gamma P^p), \quad Q(\gamma, \hat{P}^p) \rightarrow Q(\gamma, P^p), \quad K((1-\gamma)Q(\gamma, \hat{P}^p)) \rightarrow K((1-\gamma)Q(\gamma, P^p)), \quad P^u \rightarrow \bar{P}^u.$$

The first term in (16) can be bounded by  $\|h\|_\infty \sum_{j=1}^N |\alpha_j^u - \bar{\alpha}_j^u|$ . Let  $w_j^u$  and  $\bar{w}_j^u$  denote the unnormalized weights in (8), i.e.,

$$w_j^u = \varphi(y; H\nu_j^{u,\gamma}, HQ(\gamma, \hat{P}^p)H') + \frac{1}{1-\gamma}R).$$

Observe that

$$w_j^u \leq \varphi(0; 0, HQ(\gamma, \hat{P}^p)H') + \frac{1}{1-\gamma}R \leq \varphi(0; 0, \frac{1}{1-\gamma}R),$$

where the last inequality follows from  $\det(M+N) \geq \det(M)$  for arbitrary positive definite  $M$  and positive semi-definite  $N$ . The same bound holds true for  $\bar{w}_j^u$ . Notice that

$$\sum_{j=1}^N |\alpha_j^u - \bar{\alpha}_j^u| \leq \frac{1}{\text{ave } w^u} \frac{1}{N} \sum_{j=1}^N |w_j^u - \bar{w}_j^u| + \frac{1}{N} \sum_{j=1}^N \bar{w}_j^u \left| \frac{1}{\text{ave } w^u} - \frac{1}{\text{ave } \bar{w}^u} \right|$$



and  $|\text{ave } w^u - \text{ave } \bar{w}^u| \leq \frac{1}{N} \sum_{j=1}^N |w_j^u - \bar{w}_j^u|$ , where  $\text{ave } w^u = \frac{1}{N} \sum_{j=1}^N w_j^u$  and  $\text{ave } \bar{w}^u = \frac{1}{N} \sum_{j=1}^N \bar{w}_j^u$ . The  $\bar{w}_j^u$  are iid, hence,  $\text{ave } \bar{w}^u$  converges almost surely, and we conclude that  $\sum_{j=1}^N |\alpha_j^u - \bar{\alpha}_j^u| \rightarrow 0$  if

$$(18) \quad \frac{1}{N} \sum_{j=1}^N |w_j^u - \bar{w}_j^u| \rightarrow 0.$$

To show (18), we fix a compact set  $D \subset \mathbb{R}^q$ . Then we have

$$(19) \quad \begin{aligned} \frac{1}{N} \sum_{j=1}^N |w_j^u - \bar{w}_j^u| &\leq \frac{1}{N} \sum_{j=1}^N |w_j^u - \bar{w}_j^u| 1_{x_j^p \notin D} + \max_{1 \leq j \leq N: x_j^p \in D} |w_j^u - \bar{w}_j^u| \\ &\leq 2\varphi(0; 0, \frac{1}{1-\gamma}R) \frac{1}{N} \sum_{j=1}^N 1_{x_j^p \notin D} + \max_{1 \leq j \leq N: x_j^p \in D} |w_j^u - \bar{w}_j^u| \\ &\rightarrow 2\varphi(0; 0, \frac{1}{1-\gamma}R) \cdot \mathbb{P}[x_1^p \notin D], \end{aligned}$$

where the “max”-term goes to zero for reasons of uniform continuity in combination with (17).

Letting  $D \uparrow \mathbb{R}^q$ , establishes (18) by virtue of dominated convergence.

We now analyze the second term in (16). Again, we fix a compact set  $D \subset \mathbb{R}^q$ . Then we have

$$\begin{aligned} \sum_{j=1}^N \bar{\alpha}_j^u \left| \mathcal{N}(\mu_j^u, P^u)h - \mathcal{N}(\bar{\mu}_j^u, \bar{P}^u)h \right| &\leq \max_{1 \leq j \leq N: x_j^p \in D} \|h\|_\infty \int_{S_h} \left| \varphi(z; \mu_j^u, P^u) - \varphi(z; \bar{\mu}_j^u, \bar{P}^u) \right| dz \\ &\quad + 2\|h\|_\infty \sum_{j=1}^N \bar{\alpha}_j^u 1_{x_j^p \notin D} \\ &\rightarrow 2\|h\|_\infty \frac{\mathbf{E}[\bar{w}_1^u 1_{x_1^p \notin D}]}{\mathbf{E}[\bar{w}_1^u]}, \end{aligned}$$

where the “max”-term goes to zero for reasons of uniform continuity in combination with (17).

Letting  $D \uparrow \mathbb{R}^q$ , shows that also the second term in (16) converges to 0, which completes the proof.

**Proof of Lemma 1.** We set

$$Z = \exp\left(-\frac{1}{2}(x_j^p - \mu^p)' C_\gamma (x_j^p - \mu^p) + d'_\gamma (x_j^p - \mu^p)\right).$$

Then by definition

$$\text{Var}[\tilde{\alpha}_j^{u,\gamma}] = \frac{1}{N^2} \left( \frac{\mathbf{E}[Z^2]}{\mathbf{E}[Z]^2} - 1 \right).$$

The lemma follows by completing the square and using that Gaussian densities integrate to one. More precisely, for any  $x$  and any positive definite matrix  $\Gamma$ :

$$x'(C_\gamma + \Gamma)x - 2d'_\gamma x = (x - (C_\gamma + \Gamma)^{-1}d_\gamma)'(C_\gamma + \Gamma)(x - (C_\gamma + \Gamma)^{-1}d_\gamma) - d'_\gamma(C_\gamma + \Gamma)^{-1}d_\gamma.$$

Therefore

$$\mathbf{E}[Z] = (\det(P^p) \det(C_\gamma + (P^p)^{-1}))^{-1/2} \exp\left(\frac{1}{2}d'_\gamma(C_\gamma + (P^p)^{-1})^{-1}d_\gamma\right)$$

and

$$\mathbf{E}[Z^2] = (\det(P^p) \det(2C_\gamma + (P^p)^{-1}))^{-1/2} \exp(2d'_\gamma(2C_\gamma + (P^p)^{-1})^{-1}d_\gamma).$$

Taking these results together, the first claim follows.

For the second claim, we note that as  $\gamma \uparrow 1$

$$\begin{aligned} C_\gamma &\sim (1 - \gamma)H'(I - K_1'H')R^{-1}(I - HK_1)H, \\ d_\gamma &\sim (1 - \gamma)H'(I - K_1'H')R^{-1}(I - HK_1)(y - H\mu^p) \end{aligned}$$

because  $K_\gamma$  and  $Q_\gamma$  are continuous. The result then follows by a straightforward computation.

#### REFERENCES

- Alspach, D. L. and Sorenson, H. W. (1972). Nonlinear Bayesian estimation using Gaussian sum approximations. *IEEE Transactions on Automatic Control*, 17:439–448.
- Bengtsson, T., Snyder, C., and Nychka, D. (2003). Toward a nonlinear ensemble filter for high-dimensional systems. *Journal of Geophysical Research*, 108:8775.
- Beskos, A., Crisan, D., and Jasra, A. (2012). On the stability of sequential Monte Carlo methods in high dimensions. Preprint.
- Bocquet, M., Pires, C. A., and Wu, L. (2010). Beyond Gaussian statistical modeling in geophysical data assimilation. *Monthly Weather Review*, 138:2997–3023.
- Burgers, G., van Leeuwen, P. J., and Evensen, G. (1998). Analysis scheme in the ensemble Kalman filter. *Monthly Weather Review*, 126:1719–1724.
- Doucet, A., Godsill, S., and Andrieu, C. (2000). On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing*, 10:197–208.
- Dovera, L. and Della Rossa, E. (2011). Multimodal ensemble Kalman filtering using Gaussian mixture models. *Computational Geosciences*, 15:307–323.
- Drazin, P. G. and Johnson, R. S. (1989). *Solitons: An Introduction*. Cambridge University Press.
- Evensen, G. (1994). Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research*, 99:10143–10162.
- Evensen, G. (2007). *Data Assimilation*. Springer.
- Frei, M. and Künsch, H. R. (2011). Mixture ensemble Kalman filters. *Computational Statistics and Data Analysis*. In press.
- Furrer, R. and Bengtsson, T. (2007). Estimation of high-dimensional prior and posterior covariance matrices in Kalman filter variants. *Journal of Multivariate Analysis*, 98:227–255.
- Gaspari, G. and Cohn, S. E. (1999). Construction of correlation functions in two and three dimensions. *Quarterly Journal of the Royal Meteorological Society*, 125(554):723–757.
- Gneiting, T., Balabdaoui, F., and Raftery, A. E. (2007). Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society: Series B*, 69:243–268.
- Gordon, N. J., Salmond, D. J., and Smith, A. F. M. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings*, 140:107–113.
- Hoteit, I., Luo, X., and Pham, D.-T. (2012). Particle Kalman filtering: A nonlinear Bayesian framework for ensemble Kalman filters. *Monthly Weather Review*, 140:528–542.
- Houtekamer, P. L. and Mitchell, H. L. (1998). Data assimilation using an ensemble Kalman filter technique. *Monthly Weather Review*, 126:796–811.

- Houtekamer, P. L. and Mitchell, H. L. (2001). A sequential ensemble Kalman filter for atmospheric data assimilation. *Monthly Weather Review*, 129(1):123–137.
- Künsch, H. R. (2005). Recursive Monte Carlo filters: algorithms and theoretical analysis. *Annals of Statistics*, 33(5):1983–2021.
- Lawson, W. G. and Hansen, J. A. (2005). Alignment error models and ensemble-based data assimilation. *Monthly Weather Review*, 133:1687–1709.
- Le Gland, F., Monbet, V., and Tran, V. (2010). Large sample asymptotics for the ensemble Kalman filter. In Crisan, D. and Rozovskii, B., editors, *The Oxford Handbook of Nonlinear Filtering*, pages 598–634. Oxford University Press.
- Lei, J. and Bickel, P. (2012). A moment matching ensemble filter for nonlinear non-gaussian data assimilation. *Monthly Weather Review*, 139:3964–3973.
- Liu, J. S. (1996). Metropolized independent sampling with comparisons to rejection sampling and importance sampling. *Statistics and Computing*, 6(2):113–119.
- Lorenz, E. N. and Emanuel, K. A. (1998). Optimal sites for supplementary weather observations: Simulations with a small model. *Journal of the Atmospheric Sciences*, 55:399–414.
- Mandel, J. and Beezley, J. D. (2009). An ensemble Kalman-particle predictor-corrector filter for non-Gaussian data assimilation. In *Computational Science – ICCS 2009*, volume 5545 of *Lecture Notes in Computer Science*, pages 470–478. Springer.
- Musso, C., Oudjane, N., and Le Gland, F. (2001). Improving regularised particle filters. In Doucet, A., de Freitas, N., and Gordon, N., editors, *Sequential Monte Carlo Methods in Practice*, pages 247–271. Springer.
- Papadakis, N., Mémin, E., Cuzol, A., and Gengembre, N. (2010). Data assimilation with the weighted ensemble Kalman filter. *Tellus A*, 62:673–697.
- Pitt, M. K. and Shephard, N. (1999). Filtering via simulation: auxiliary particle filters. *Journal of the American Statistical Association*, 94(446):590–599.
- Randles, R. H. (1982). On the asymptotic normality of statistics with estimated parameters. *Annals of Statistics*, 10:462–474.
- Rezaie, J. and Eidsvik, J. (2012). Shrunked  $(1 - \alpha)$  ensemble Kalman filter and  $\alpha$  Gaussian mixture filter. *Computational Geosciences*, 16:837–852.
- Sakov, P. and Bertino, L. (2010). Relation between two common localisation methods for the EnKF. *Computational Geosciences*, 15(2):225–237.
- Snyder, C., Bengtsson, T., Bickel, P., and Anderson, J. (2008). Obstacles to high-dimensional particle filtering. *Monthly Weather Review*, 136:4629–4640.
- Sun, A. Y., Morris, A. P., and Mohanty, S. (2009). Sequential updating of multimodal hydrogeologic parameter fields using localization and clustering techniques. *Water Resources Research*, 45:W07424.
- van Leeuwen, P. J. (2003). A variance-minimizing filter for largescale applications. *Monthly Weather Review*, 131:2071–2084.
- Zupanski, D. and Zupanski, M. (2006). Model error estimation employing an ensemble data assimilation approach. *Monthly Weather Review*, 134:1337–1354.